

## CLAIMS

1. A network system, comprising:  
a first network node, comprising:  
an input for receiving packets, wherein during operation the first network  
node operates as a downstream node when receiving packets at the input from an  
5 upstream node;  
a buffer, coupled to the input and for storing received packets;  
circuitry for detecting when a number of packets stored in the buffer  
exceeds a buffer storage threshold; and  
circuitry, responsive to a detection by the circuitry for detecting that the  
10 number of packets stored in the buffer exceeds the buffer storage threshold, for issuing a  
message to the upstream node, wherein the message selectively commands the upstream  
node to reduce a rate of transmission of packets from the upstream node to the  
downstream node to a non-zero rate.
2. The system of claim 1:  
wherein the message comprises a first message;  
wherein the buffer storage threshold comprises a first buffer storage threshold;  
wherein the circuitry for detecting further comprising circuitry for detecting when  
5 a number of packets stored in the buffer exceeds a second buffer storage threshold,  
wherein the second buffer storage threshold is larger than the first buffer storage  
threshold;  
wherein the circuitry for issuing is for, responsive to a detection by the circuitry for  
detecting that the number of packets stored in the buffer exceeds the second buffer storage  
10 threshold, issuing a second message to the upstream node; and  
wherein the second message commands the upstream node to cease transmission  
of packets from the upstream node to the downstream node.

3. The system of claim 2:

wherein the second message commands the upstream node to cease transmission of packets from the upstream node to the downstream node for a period of time; and

wherein the second message identifies the period of time.

4. The system of claim 3 wherein the first network node further comprises circuitry for determining the period of time in response to a rate of transmission from the buffer and a measure of packet occupancy in the buffer.

5. The system of claim 1 wherein the message identifies the non-zero rate.

6. The system of claim 5 wherein the node further comprises circuitry for determining the non-zero rate by multiplying a constant, less than one, times a rate that packets are output from the buffer at a time that the circuitry for detecting detects that the number of packets stored in the buffer exceeds the buffer storage threshold.

7. The system of claim 1:

wherein the message comprises a first message;

wherein the buffer storage threshold comprises a first buffer storage threshold;

5 wherein the circuitry for detecting further comprises circuitry for detecting when a number of packets stored in the buffer exceeds a second buffer storage threshold, wherein the second buffer storage threshold is larger than the first buffer storage threshold;

wherein the circuitry for issuing is for, responsive to a detection by the circuitry for detecting that the number of packets stored in the buffer exceeds the second buffer storage threshold, issuing a second message to the upstream node;

10 wherein the first message selectively commands the upstream node to either continue with transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence in the buffer of the downstream node relative to a packet drop precedence in at least a portion of a buffer in  
15 the upstream node; and

wherein the second message selectively commands the upstream node to either cease transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence in the buffer of the  
20 downstream node relative to a packet drop precedence in the at least a portion of the buffer in the upstream node.

8. The system of claim 7 wherein the at least a portion of the buffer in the upstream node comprises a number of locations in the buffer storing packets to be transmitted at the non-zero rate and during a period of time identified by the message.

9. The system of claim 7 wherein the at least a portion of the buffer in the upstream node comprises a number of locations in the buffer less than the buffer entirety.

10. The system of claim 7:

wherein the first message selectively commands the upstream node to either continue with transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence dominance in the buffer of the downstream node relative to a packet drop precedence dominance in at least a portion of a buffer in the upstream node; and

wherein the second message selectively commands the upstream node to either cease transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence dominance in the buffer of the downstream node relative to a packet drop precedence dominance in the at least a portion of the buffer in the upstream node.

11. The system of claim 10 wherein the packet drop precedence of the downstream node and of the upstream node comprises a color marked precedence.

12. The system of claim 11:

wherein the packet drop precedence of the downstream node and of the upstream node ranks packet priority, from highest to lowest, as green, yellow, and red;

5 wherein responsive to a red packet drop precedence dominance in the buffer in the downstream node and a red packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate;

10 wherein responsive to a red packet drop precedence dominance in the buffer in the downstream node and a drop precedence dominance other than red in the in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to continue a present rate of transmission of packets from the upstream node to the downstream node;

15 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and packets having either a red, yellow, or combined red and yellow, packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate;

20 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and a drop precedence dominance in the at least a portion of the buffer in the upstream node other than red, yellow, or combined red and yellow, the first message commands the upstream node to continue a present rate of transmission of packets from the upstream node to the downstream node; and

25 wherein responsive to a green packet drop precedence dominance in the buffer of the downstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate.

13. The system of claim 11:

wherein the packet drop precedence of the downstream node and of the upstream node ranks packet priority, from highest to lowest, as green, yellow, and red;

5 wherein responsive to a red packet drop precedence dominance in the buffer of the downstream node and a red packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the second message commands the upstream node to cease transmission of packets from the upstream node to the downstream node;

10 wherein responsive to a red packet drop precedence dominance in the buffer of the downstream node and a drop precedence dominance other than red in the at least a portion of the buffer in the upstream node, the second message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate;

15 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and packets having either a red, yellow, or combined red and yellow, packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the second message commands the upstream node to cease transmission of packets from the upstream node to the downstream node;

20 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and a drop precedence dominance in the at least a portion of the buffer in the upstream node other than red, yellow, or combined red and yellow, the second message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate; and

25 wherein responsive to a green packet drop precedence dominance in the buffer of the downstream node, the second message commands the upstream node to cease transmission of packets from the upstream node to the downstream node.

14. The system of claim 13:

wherein the packet drop precedence of the downstream node and of the upstream node ranks packet priority, from highest to lowest, as green, yellow, and red;

5 wherein responsive to a red packet drop precedence dominance in the buffer of the downstream node and a red packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate;

10 wherein responsive to a red packet drop precedence dominance in the buffer of the downstream node and a drop precedence dominance other than red in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to continue a present rate of transmission of packets from the upstream node to the downstream node;

15 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and packets having either a red, yellow, or combined red and yellow, packet drop precedence dominance in the at least a portion of the buffer in the upstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate;

20 wherein responsive to a yellow packet drop precedence dominance in the buffer of the downstream node and a drop precedence dominance in the at least a portion of the buffer in the upstream node other than red, yellow, or combined red and yellow, the first message commands the upstream node to continue a present rate of transmission of packets from the upstream node to the downstream node; and

25 wherein responsive to a green packet drop precedence dominance in the buffer of the downstream node, the first message commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate.

15. The system of claim 7 wherein the packet drop precedence of the downstream node and of the upstream node comprises a color marked precedence.

16. The system of claim 1 wherein the upstream node is adjacent the downstream node.

17. The system of claim 1 wherein the system further comprises the upstream node.

18. The system of claim 1:

wherein the message selectively commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate for a period of time; and

5 wherein the first network node further comprises circuitry for determining the period of time.

19. The system of claim 18:

wherein the circuitry for determining the period of time determines the period of time in response to a transmission rate of packets from the buffer and a receipt rate of packets into the buffer; and

5 wherein the transmission rate is greater than the receipt rate.



20. A method of operating a first network node, comprising:

receiving packets at an input of the first network node, wherein during operation the first network node operates as a downstream node when receiving packets at the input from an upstream node;

5 storing the received packets into a buffer;

detecting when a number of packets stored in the buffer exceeds a buffer storage threshold; and

responsive to a detecting that the number of packets stored in the buffer exceeds the buffer storage threshold, issuing a message to the upstream node, wherein the  
10 message selectively commands the upstream node to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate.

21. The method of claim 20:

wherein the message comprises a first message;

wherein the buffer storage threshold comprises a first buffer storage threshold;

wherein the step of detecting is further for detecting when a number of packets  
5 stored in the buffer exceeds a second buffer storage threshold, wherein the second buffer storage threshold is larger than the first buffer storage threshold;

and further comprising, responsive to detecting that the number of packets stored in the buffer exceeds the second buffer storage threshold, issuing a second message to the  
upstream node; and

10 wherein the second message commands the upstream node to cease transmission of packets from the upstream node to the downstream node.

22. The method of claim 20:

wherein the message comprises a first message;

wherein the buffer storage threshold comprises a first buffer storage threshold;

5 wherein the step of detecting further comprises detecting when a number of packets stored in the buffer exceeds a second buffer storage threshold, wherein the second buffer storage threshold is larger than the first buffer storage threshold;

wherein the step of issuing is for, responsive to a detection that the number of packets stored in the buffer exceeds the second buffer storage threshold, issuing a second message to the upstream node;

10 wherein the first message selectively commands the upstream node to either continue with transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence in the buffer of the downstream node relative to a packet drop precedence in at least a portion of a buffer in  
15 the upstream node; and

wherein the second message selectively commands the upstream node to either cease transmission of packets from the upstream node to the downstream node or to reduce a rate of transmission of packets from the upstream node to the downstream node to a non-zero rate, in response to a packet drop precedence in the buffer of the  
20 downstream node relative to a packet drop precedence in the at least a portion of the buffer in the upstream node.

23. A network system, comprising:  
a first network node, comprising:  
a buffer for storing packets;  
circuitry for transmitting packets from the buffer, as an upstream node, to a  
5 downstream node; and  
an input for receiving a message from the downstream node, wherein the  
circuitry for transmitting selectively reduces its rate of transmission of packets from the  
upstream node to the downstream node to a non-zero rate in response to the message.

24. The system of claim 23 when the message is issued by the downstream  
node to the upstream node in response to the downstream node detecting that a number  
of packets stored in a buffer in the downstream node exceeds a buffer storage threshold  
associated with the buffer in the downstream node.

25. The system of claim 24:  
wherein the message comprises a first message;  
wherein the buffer storage threshold comprises a first buffer storage threshold;  
wherein the input is further for receiving a second message from the downstream  
5 node, wherein the circuitry for transmitting selectively ceases its rate of transmission of  
packets from the upstream node to the downstream node in response to the second  
message.

26. The system of claim 25 wherein the wherein the circuitry for transmitting  
selectively ceases its rate of transmission of packets from the upstream node to the  
downstream node in response to the second message and in response to a packet drop  
precedence in at least a portion of the buffer of the upstream node relative to a packet drop  
5 precedence in a buffer in the downstream node.

27. The system of claim 25 wherein the second message is issued by the downstream node when a number of packets stored in the buffer of the downstream node exceeds a second buffer storage threshold, wherein the second buffer storage threshold is larger than the first buffer storage threshold.

28. The system of claim 23 wherein the wherein the circuitry for transmitting selectively reduces its rate of transmission of packets from the upstream node to the downstream node to a non-zero rate in response to the message and in response to a packet drop precedence in at least a portion of the buffer of the upstream node relative to a  
5 packet drop precedence in a buffer in the downstream node.

29. The system of claim 28 wherein the packet drop precedence of the downstream node and of the upstream node comprises a color marked precedence.

30. A method of operating a first network node, comprising:  
storing packets in a buffer;  
transmitting packets from the buffer, as an upstream node, to a  
downstream node; and

5 receiving a message from the downstream node, wherein the step of  
transmitting selectively reduces its rate of transmission of packets from the upstream node  
to the downstream node to a non-zero rate in response to the message.

31. The method of claim 30 when the message is issued by the downstream  
node to the upstream node in response to the downstream node detecting that a number  
of packets stored in a buffer in the downstream node exceeds a buffer storage threshold  
associated with the buffer in the downstream node.

32. The method of claim 31:  
wherein the message comprises a first message;  
wherein the buffer storage threshold comprises a first buffer storage threshold;  
and further comprising receiving a second message from the downstream node,  
5 wherein the step of transmitting ceases its transmission of packets from the upstream  
node to the downstream node in response to the second message.

33. The method of claim 32 wherein the second message is issued by the  
downstream node when a number of packets stored in the buffer of the downstream node  
exceeds a second buffer storage threshold, wherein the second buffer storage threshold is  
larger than the first buffer storage threshold.

34. The method of claim 32 wherein the step of transmitting ceases its  
transmission of packets from the upstream node to the downstream node in response to  
the second message and in response to a packet drop precedence in at least a portion of  
the buffer of the upstream node relative to a packet drop precedence in a buffer in the  
5 downstream node.

35. The method of claim 30 wherein the step of transmitting selectively reduces its rate of transmission of packets from the upstream node to the downstream node to a non-zero rate in response to the message and in response to a packet drop precedence in at least a portion of the buffer of the upstream node relative to a packet drop  
5 precedence in a buffer in the downstream node.

36. The method of claim 35 wherein the packet drop precedence of the downstream node and of the upstream node comprises a color marked precedence.

\* \* \* \* \*